

US009733843B2

(12) **United States Patent**
Callaway et al.

(10) **Patent No.:** **US 9,733,843 B2**
(45) **Date of Patent:** ***Aug. 15, 2017**

(54) **DYNAMIC CACHING MODULE SELECTION FOR OPTIMIZED DATA DEDUPLICATION**

(58) **Field of Classification Search**
CPC G06F 12/0891; G06F 12/0871
See application file for complete search history.

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

(56) **References Cited**

(72) Inventors: **Robert D. Callaway**, Holly Springs, NC (US); **Ioannis Papapanagiotou**, Raleigh, NC (US)

U.S. PATENT DOCUMENTS

8,086,799 B2 12/2011 Mondal et al.
9,298,637 B2* 3/2016 Callaway G06F 12/0891
(Continued)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

FOREIGN PATENT DOCUMENTS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

CN 102195877 A 9/2011
CN 102546817 A 7/2012

This patent is subject to a terminal disclaimer.

OTHER PUBLICATIONS

Papapanagiotou et al., "Chunk and Object Level Deduplication for Web Optimization: A Hybrid Approach," International Communications Conference (ICC), 2012 IEEE, Jun. 2012.

(21) Appl. No.: **15/068,535**

Primary Examiner — Edward Dudek, Jr.

(22) Filed: **Mar. 11, 2016**

(74) *Attorney, Agent, or Firm* — Steven M. Greenberg, Esq.; CRGO Law

(65) **Prior Publication Data**

US 2016/0224256 A1 Aug. 4, 2016

Related U.S. Application Data

(57) **ABSTRACT**

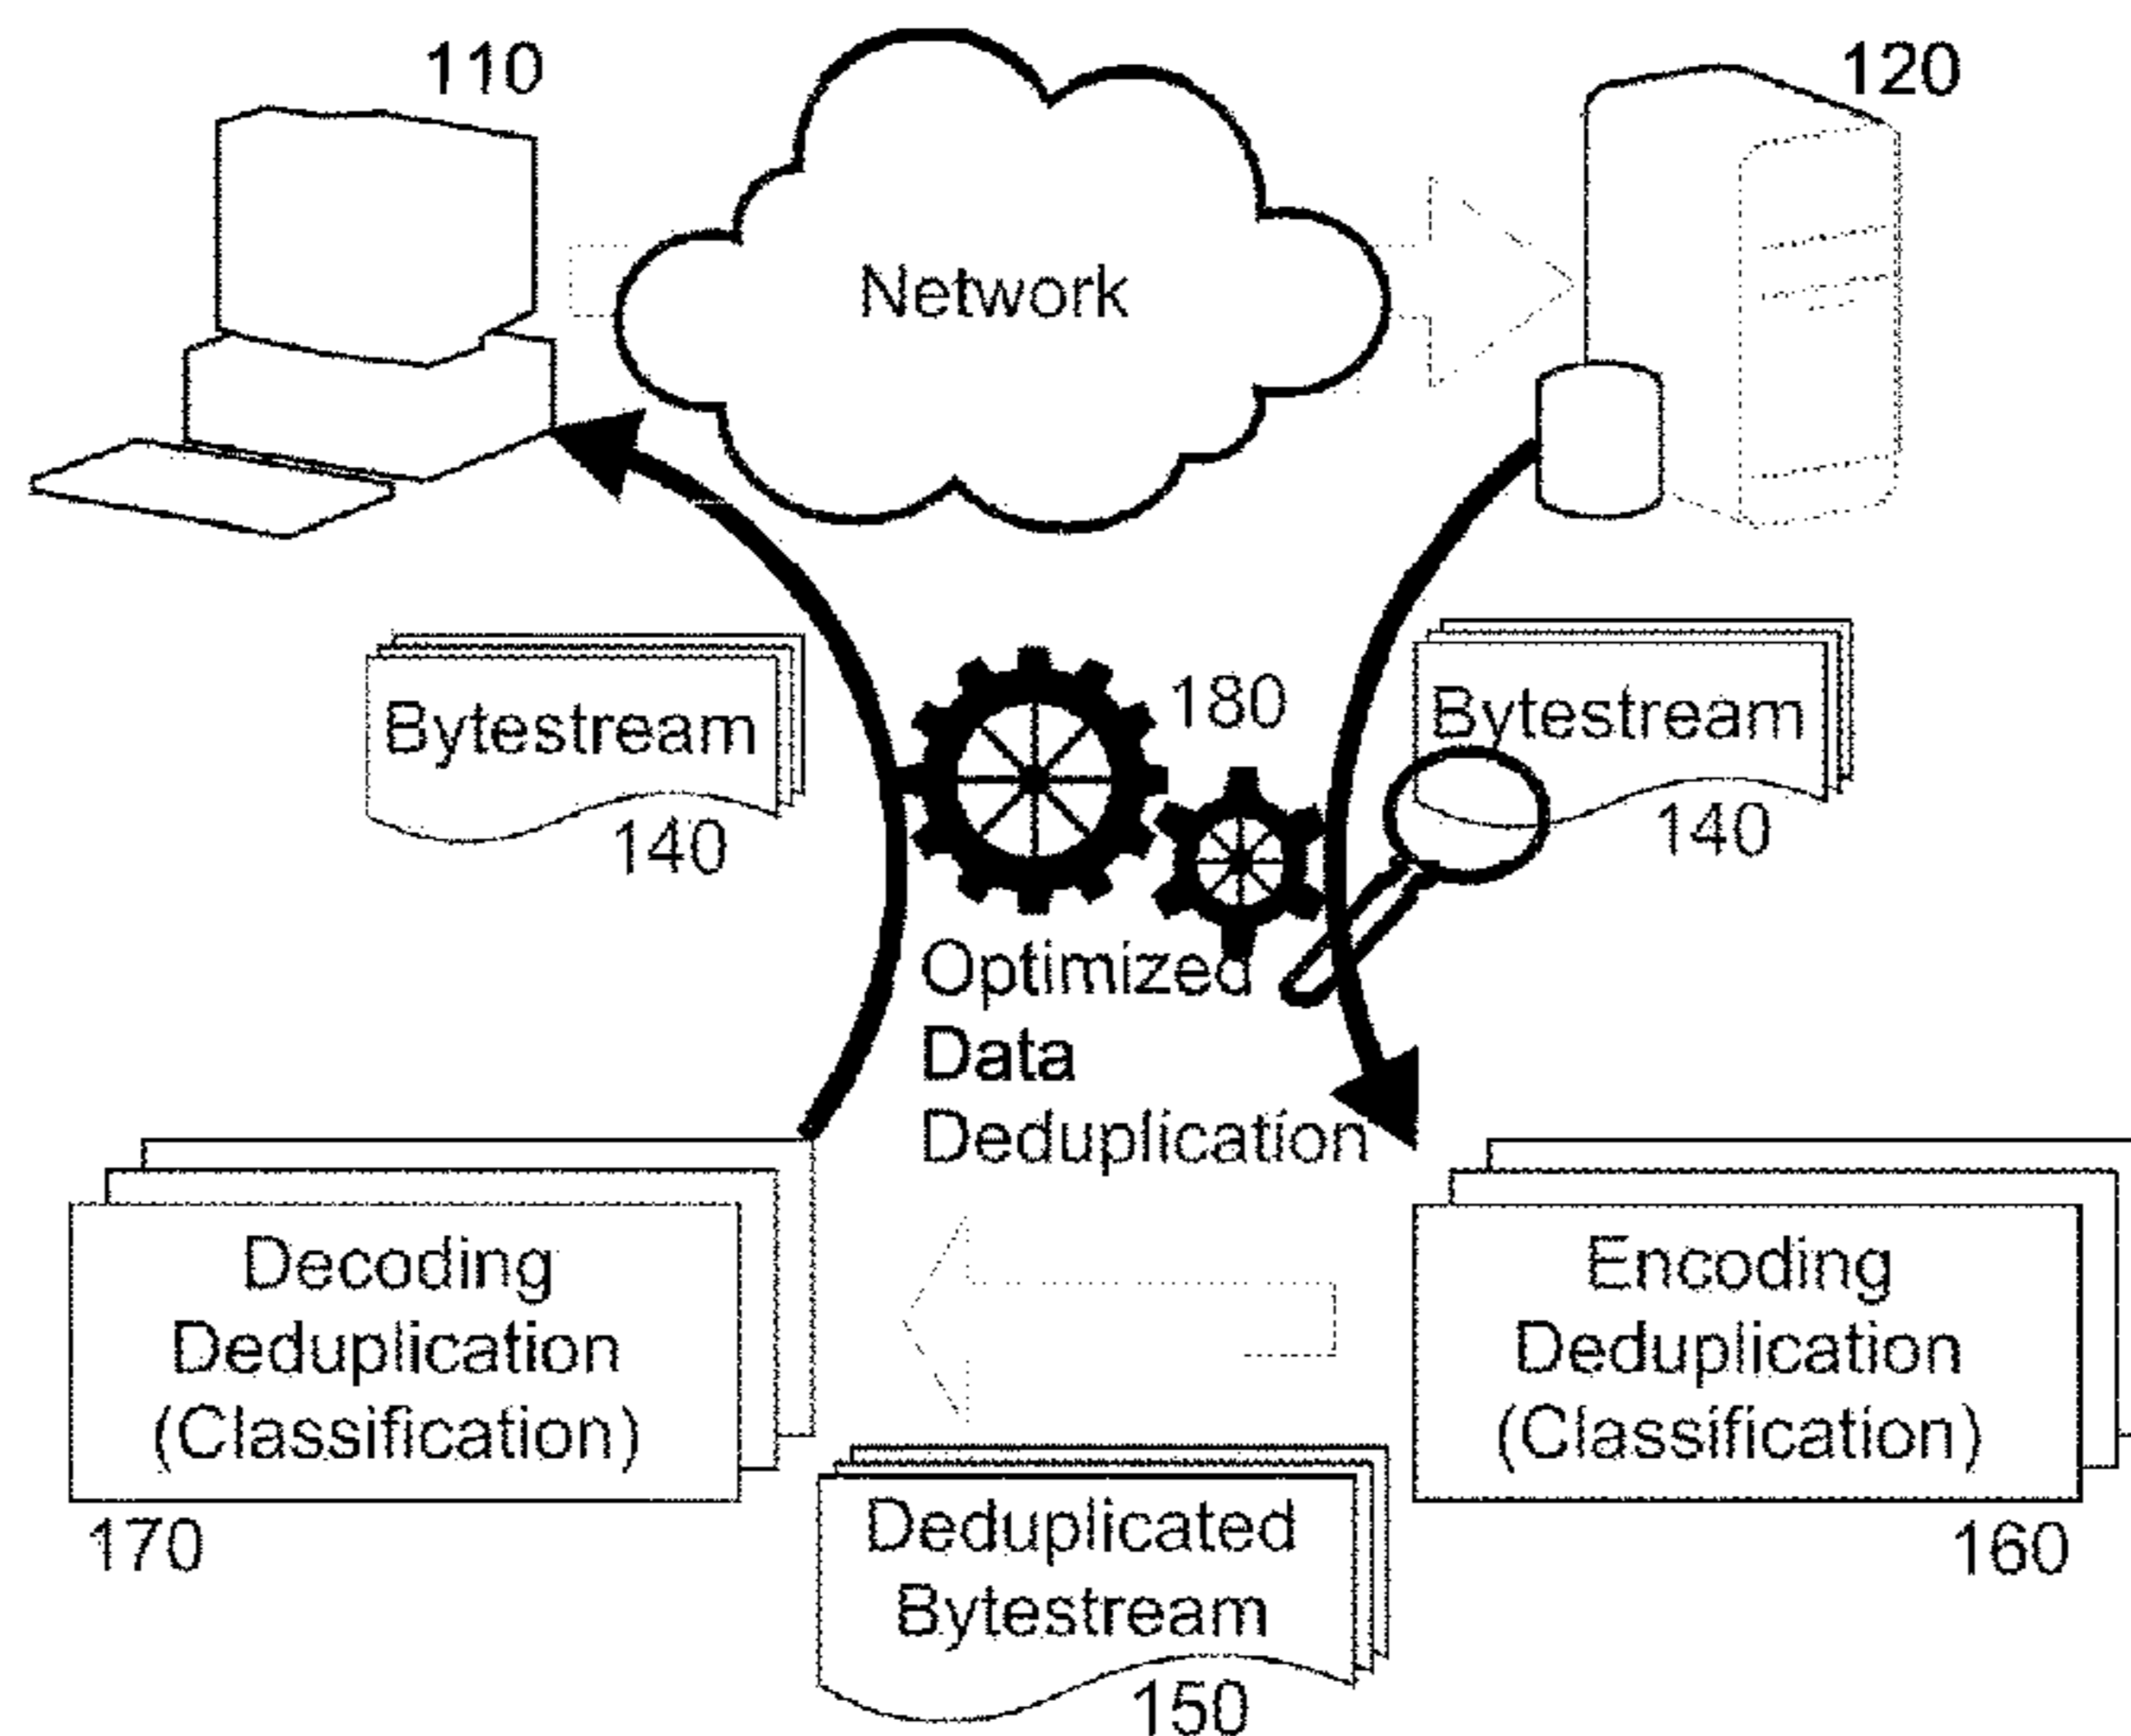
(63) Continuation of application No. 14/059,959, filed on Oct. 22, 2013, now Pat. No. 9,298,638, which is a (Continued)

Embodiments of the invention provide a method, system and computer program product for dynamic caching module selection for optimized data deduplication. In an embodiment of the invention, a method for dynamic caching module selection for optimized data deduplication is provided. The method includes receiving a request to retrieve data and classifying the request. The method also includes identifying from amongst multiple different caching modules each with a different configuration a particular caching module associated with the classification of the request. Finally, the method includes deduplicating the data in the identified caching module.

(51) **Int. Cl.**
G06F 12/08 (2016.01)
G06F 3/06 (2006.01)
(Continued)

9 Claims, 2 Drawing Sheets

(52) **U.S. Cl.**
CPC **G06F 3/0613** (2013.01); **G06F 3/067** (2013.01); **G06F 3/0641** (2013.01);
(Continued)



Related U.S. Application Data

continuation of application No. 13/800,289, filed on
Mar. 13, 2013, now Pat. No. 9,298,637.

(51) **Int. Cl.**

G06F 12/0891 (2016.01)

G06F 12/0871 (2016.01)

G06F 12/0831 (2016.01)

G06F 12/0855 (2016.01)

(52) **U.S. Cl.**

CPC **G06F 3/0653** (2013.01); **G06F 12/0833**
(2013.01); **G06F 12/0855** (2013.01); **G06F**
12/0871 (2013.01); **G06F 12/0891** (2013.01);
G06F 2212/264 (2013.01); **G06F 2212/283**
(2013.01); **G06F 2212/314** (2013.01); **G06F**
2212/621 (2013.01)

(56)

References Cited

U.S. PATENT DOCUMENTS

9,298,638 B2 *	3/2016	Callaway	G06F 12/0891
2009/0182789 A1	7/2009	Sandorfi et al.	
2011/0228674 A1 *	9/2011	Pais	H04L 49/90 370/235
2012/0047328 A1	2/2012	Williams et al.	
2013/0265869 A1 *	10/2013	Zhang	H04L 47/36 370/229
2014/0281262 A1	9/2014	Callaway et al.	711/135

* cited by examiner

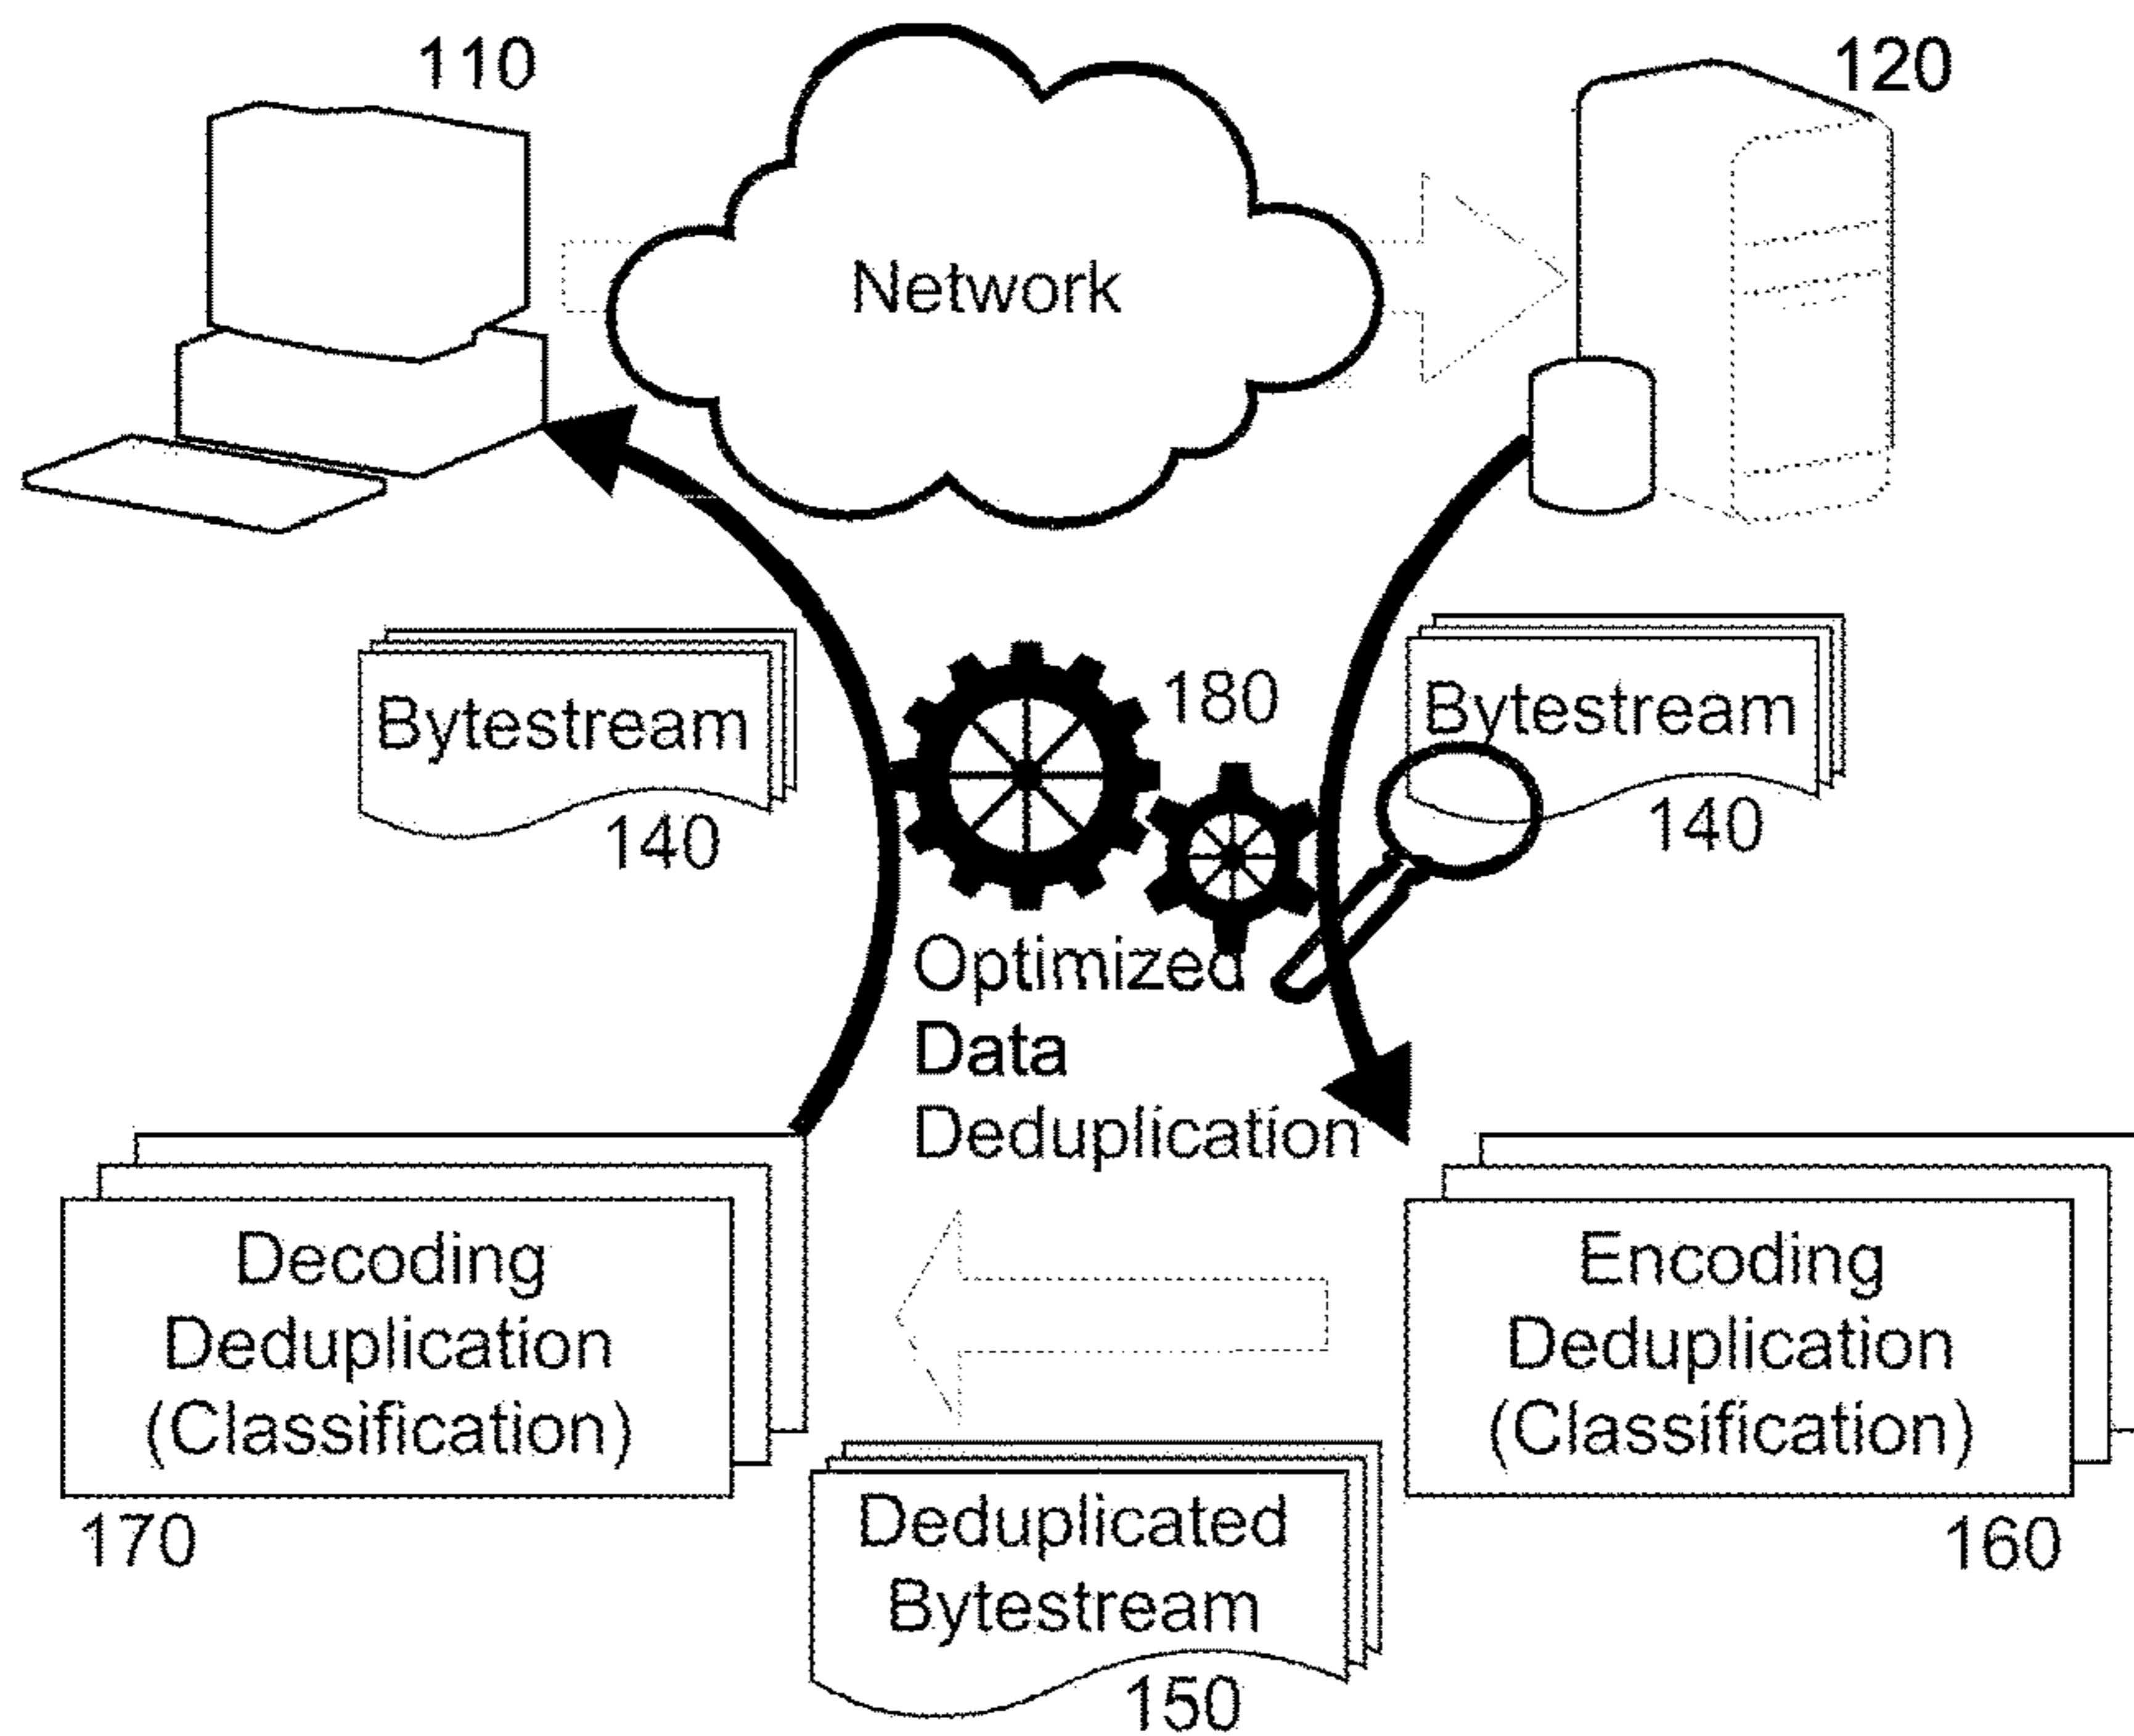


FIG. 1

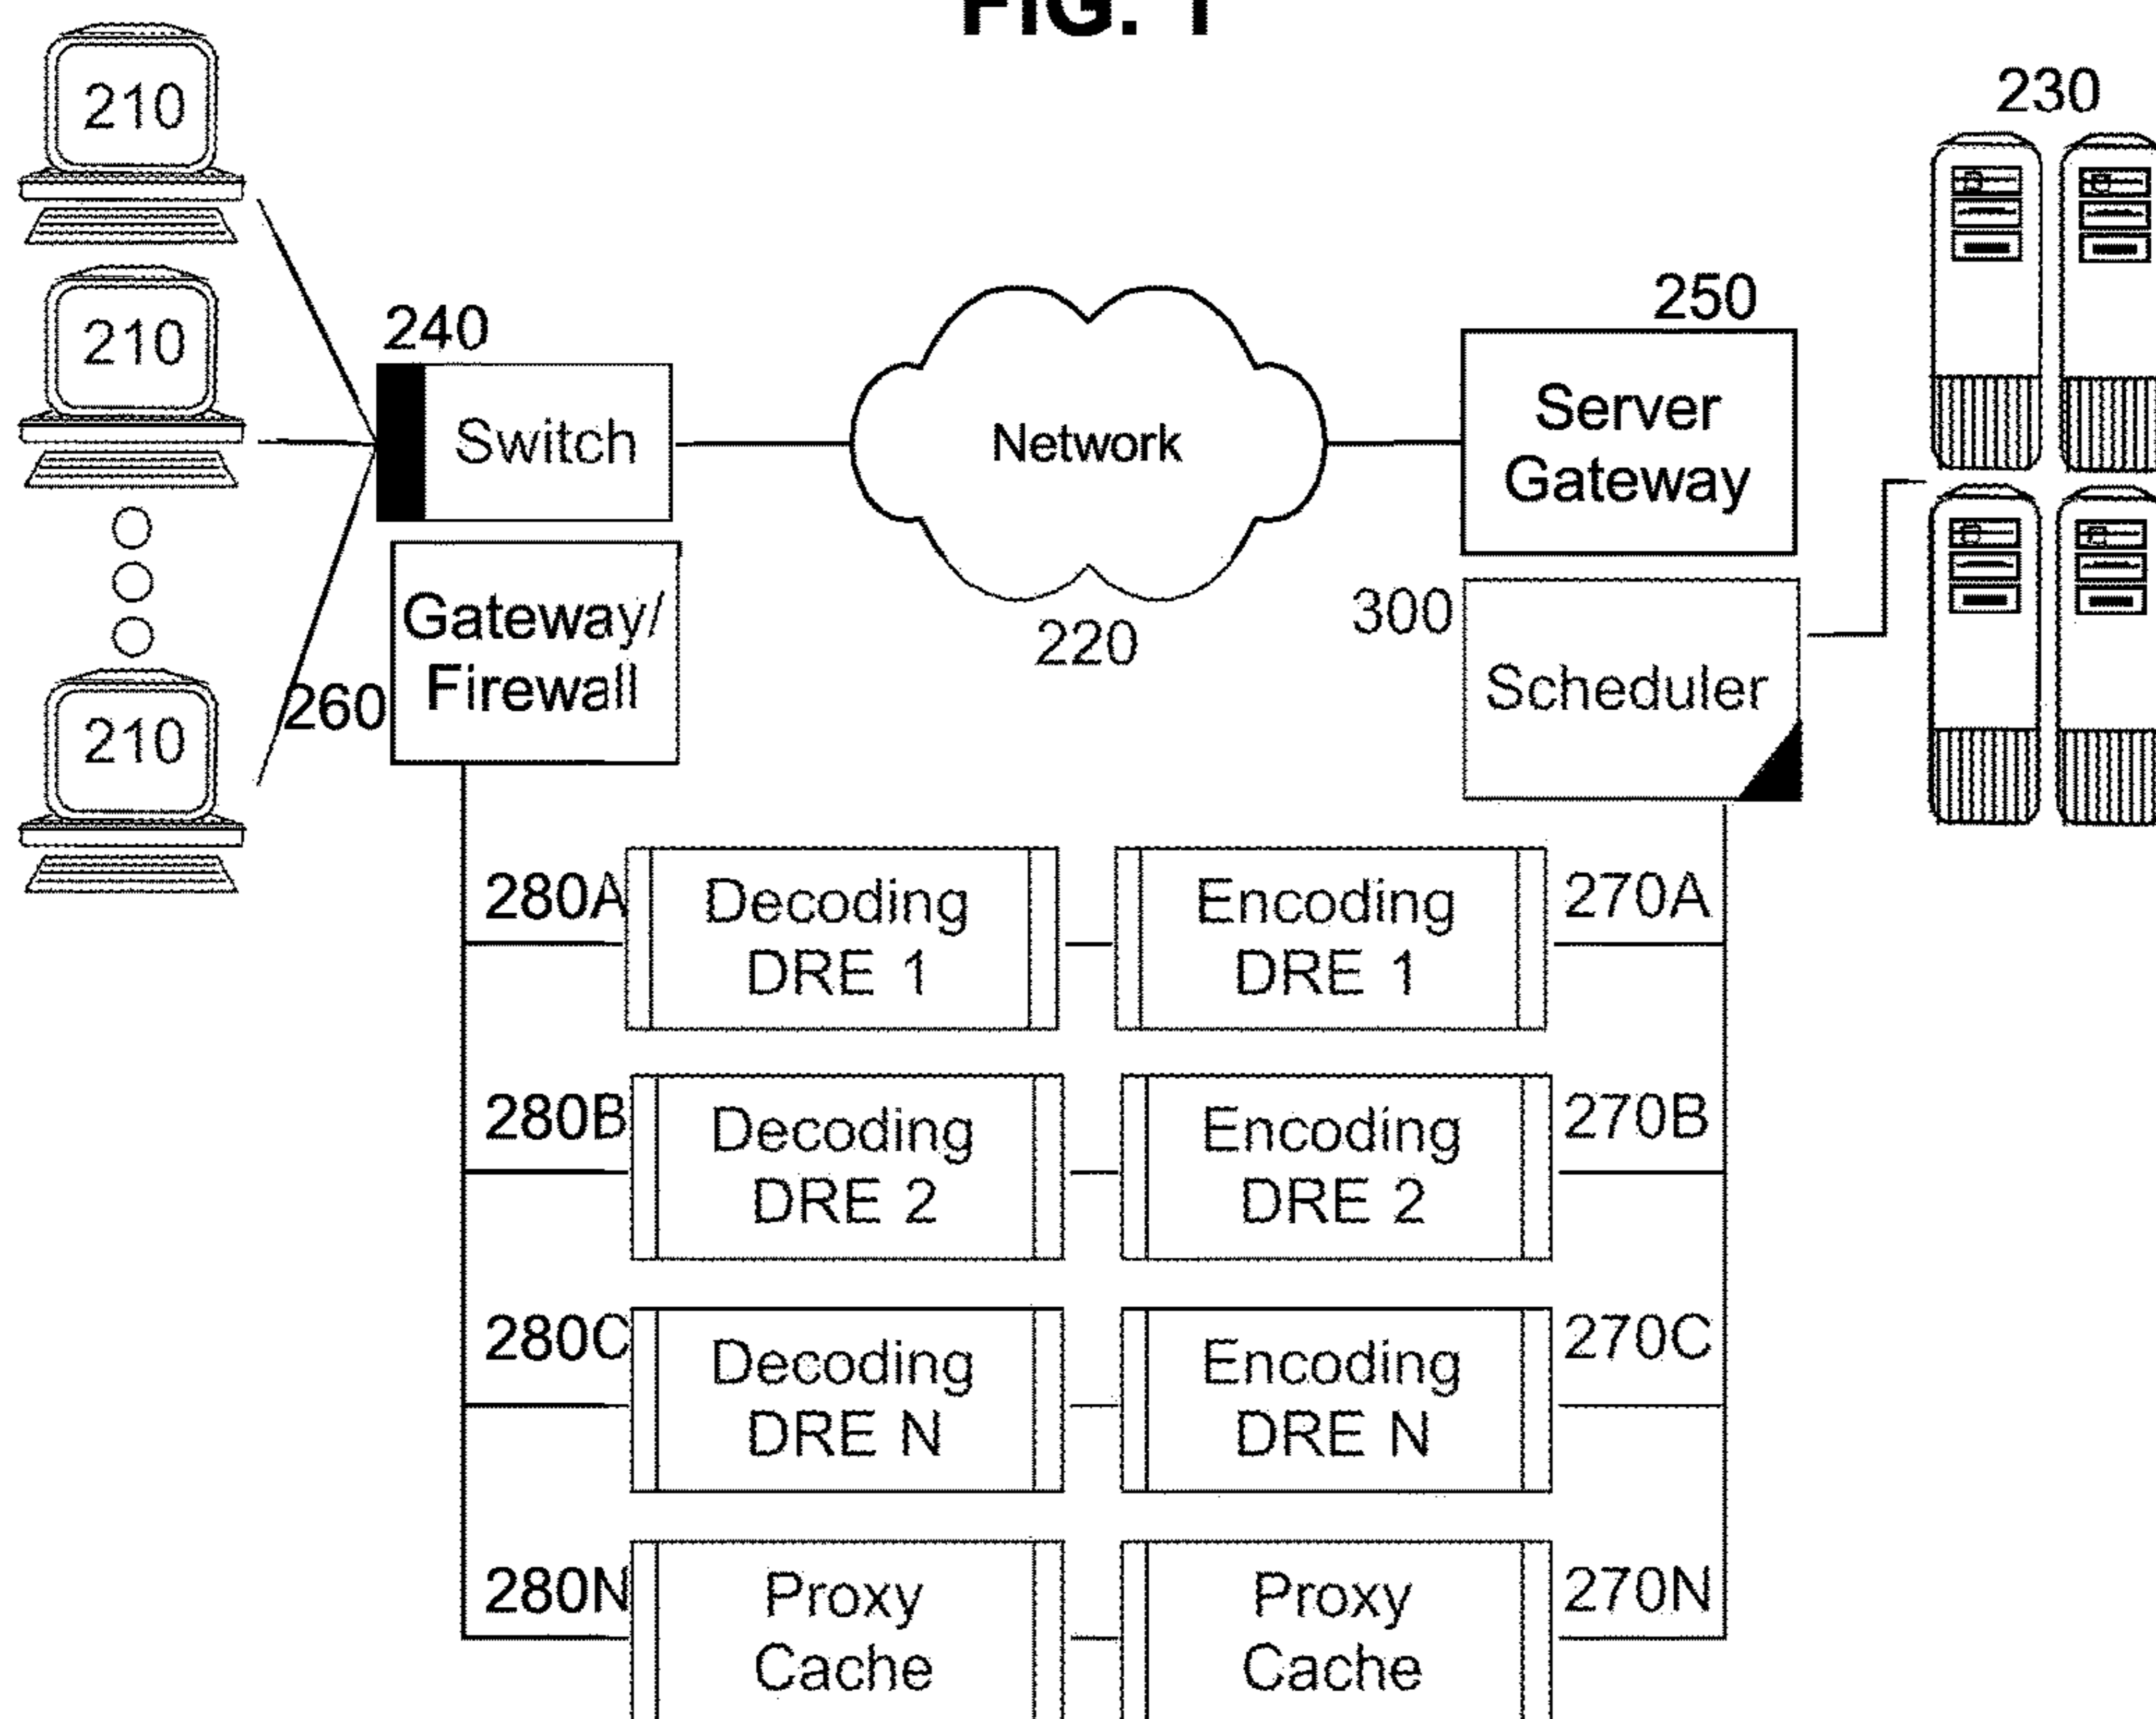


FIG. 2

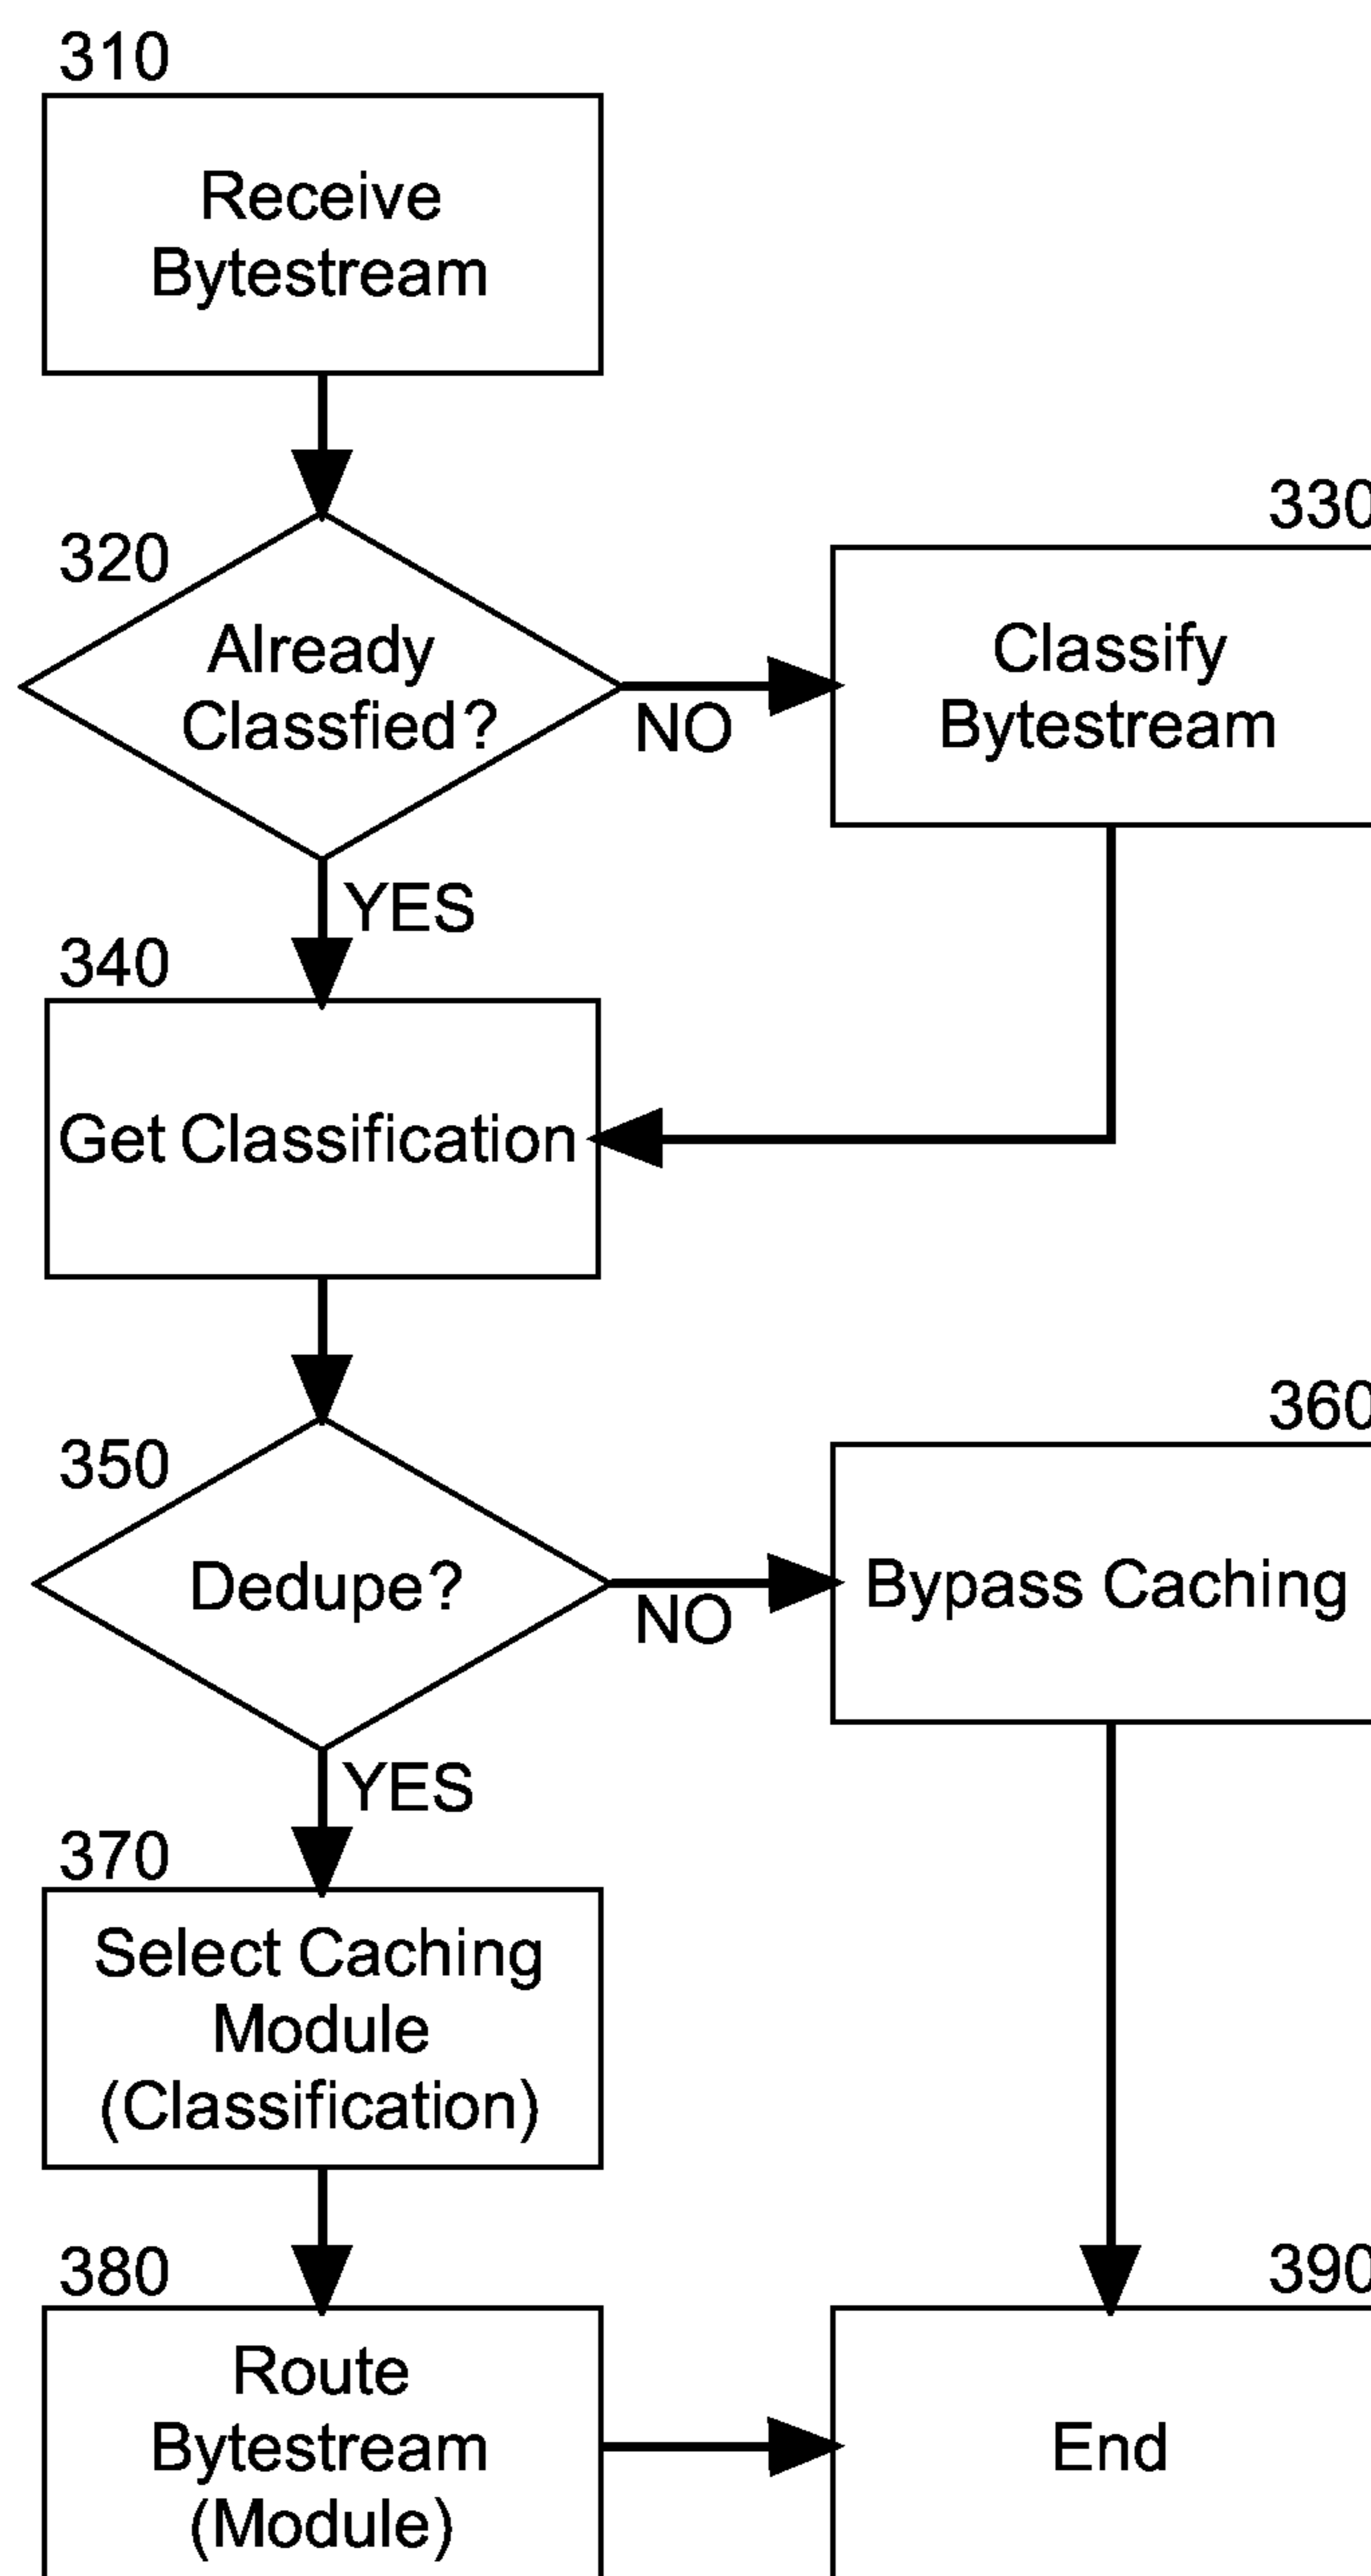


FIG. 3

DYNAMIC CACHING MODULE SELECTION FOR OPTIMIZED DATA DEDUPLICATION

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a Continuation of U.S. patent application Ser. No. 14/059,959, filed on Oct. 22, 2013, which is a Continuation of U.S. patent application Ser. No. 13/800,289, filed Mar. 13, 2013, both of which are incorporated herein by reference in their entirety.

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to data deduplication and more particularly to data deduplication through byte caching.

Description of the Related Art

Data deduplication refers to the reduction of a data flow through the elimination of redundancies in data. The act of data deduplication can reduce the quantity of traffic transmitted across a communications channel thereby increasing the responsiveness of communications between network entities exchanging data over the communications channel. Data deduplication can be performed technically in several different ways including data compression, delta encoding, proxy caching and data redundancy elimination.

Data compression removes redundant content on a per-object basis by representing duplicate bytes with hash values. In delta encoding, a technique applicable only to Web based objects, similar portions of a Web object can be represented with a hash value. Proxy caching, like data compression and delta encoding, is an object based method that performs object level deduplication by storing an object that may potentially be referenced later. Data redundancy elimination, also referred to as byte caching, differs from data compression, delta encoding and proxy caching in that byte caching operates at the byte level and is not limited only to Web objects and the hypertext transfer protocol (HTTP).

In byte caching, a combination of an encoder and decoder act in concert over a sideband channel to identify within the encoder regions of repeated bytes within byte streams and to replace these regions with hash values only to be reconstituted as a full byte stream at the decoder. The precision afforded by byte caching can provide the most effective form of deduplication, but not without substantial computational and resource cost.

Specifically, the use of byte caching as middleware in a data processing system can result in undesirable excessive memory utilization, and overutilization of processing cycles creating a bottleneck of throughput at the point of byte caching. Further, to the extent that byte caching relies upon proper fingerprint size selection which can vary in effectiveness for data from different application sources, the utilization of byte caching for data flows of different applications can be effective only for some applications and not others. Finally, much of the effectiveness of byte caching is mitigated through the inefficiency of byte caching upon byte streams lacking redundancy such as encrypted byte streams.

BRIEF SUMMARY OF THE INVENTION

Embodiments of the present invention address deficiencies of the art in respect to byte caching for data deduplication and provide a novel and non-obvious method, system and computer program product for dynamic caching module

selection for optimized data deduplication. In an embodiment of the invention, a method for dynamic caching module selection for optimized data deduplication is provided. The method includes receiving a request to retrieve data and classifying the request. The method also includes identifying from amongst multiple different caching modules each with a different configuration a particular caching module associated with the classification of the request. Finally, the method includes deduplicating the data in the identified caching module.

In one aspect of the embodiment, the request is classified according to a table correlating different requests for different ones of the caching modules, for example a table correlates a protocol and network address for each of the different requests with a corresponding one of the caching modules. The table can include entries determined by processing training data for each of the different requests in each of the caching modules and correlating each of the different requests with an optimal one of the modules. In another aspect of the embodiment, the caching modules include byte caching modules each configured with a different fingerprint size. In yet another aspect of the embodiment, the caching modules additionally include an object caching module. In even yet another aspect of the embodiment, deduplication of the data is bypassed when the request indicates that data is encrypted.

In another embodiment of the invention, a data deduplication data processing system is configured for dynamic caching module selection for optimized data deduplication. The system includes a server farm communicatively coupled to different client computers over a computer communications network and middleware disposed between the server farm and the client computers and executing in memory of a host computer. The middleware includes different caching modules, each caching module having a different configuration. Finally, the system includes a scheduler that includes program code executing in the memory of the host computer and enabled to classify a request to retrieve data from the server farm, to identify from amongst the different caching modules a particular caching module associated with the classification of the request, and to route the data for deduplication in the identified caching module.

Additional aspects of the invention will be set forth in part in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention. The aspects of the invention will be realized and attained by means of the elements and combinations particularly pointed out in the appended claims. It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention, as claimed.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute part of this specification, illustrate embodiments of the invention and together with the description, serve to explain the principles of the invention. The embodiments illustrated herein are presently preferred, it being understood, however, that the invention is not limited to the precise arrangements and instrumentalities shown, wherein:

FIG. 1 is a pictorial illustration of a process for dynamic caching module selection for optimized data deduplication;

FIG. 2 is a schematic illustration of a data deduplication data processing system configured for dynamic caching module selection for optimized data deduplication; and,

FIG. 3 is a flow chart illustrating a process for dynamic caching module selection for optimized data deduplication.

DETAILED DESCRIPTION OF THE INVENTION

Embodiments of the invention provide for dynamic caching module selection for optimized data deduplication. In accordance with an embodiment of the invention, an inbound byte stream can be classified and compared to a table of pre-determined classifications correlating specific classifications to differently configured byte stream caches with different fingerprint sizes. A particular one of the byte stream caches can be selected according to the classification of the inbound byte stream and the table and the byte stream can be routed for deduplication in the selected byte stream cache. Optionally, a proxy cache can be included in the table as well such that byte streams of a particular classification can be routed to the proxy cache for deduplication. As a further option, deduplication can be bypassed entirely for particular classified byte streams such as encrypted byte streams.

In further illustration, FIG. 1 pictorially shows a process for dynamic caching module selection for optimized data deduplication. As shown in FIG. 1, a data source 120 can respond to a data request from client 110 over network 130 by locating a requested byte stream 140. Optimized data deduplication logic 180 can classify the byte stream 140, for example in terms of a consuming application in the client 110, an address of the consuming application in the client 110, and a port in the client 110 through which the consuming application is accessed.

Thereafter, the optimized data deduplication logic 180 can select a particular one of several differently configured encoding deduplication modules 160 associated with the classification to perform byte caching on the byte stream 140. The resulting byte caching can generate a deduplicated byte stream 150 that is provided over a sideband channel to a corresponding one of several differently configured decoding deduplication modules 170. Finally, the corresponding one of the decoding deduplication modules 170 can act upon the deduplicated byte stream 150 to regenerate the byte stream 140 for delivery to the client 110.

The process described in connection with FIG. 1 can be implemented within a data deduplication data processing system. In yet further illustration, FIG. 2 schematically shows a data deduplication data processing system configured for dynamic caching module selection for optimized data deduplication. The system can include a server farm 230 including data storage that manages the storage and retrieval of data from data storage across one or more server computers and communicatively coupled to data consuming clients 210 by way of a computer communications network 220 through a server gateway 250. Of note, optimized data deduplication middleware can be disposed between the server farm 230 and the clients 210.

The middleware can include multiple different encoding caching modules 270A, 270B, 270C each executing in the memory of a host computer and communicatively linked to corresponding decoding caching modules 280A, 280B, 280C. The middleware also can include a proxy cache 270N, 280N configured to act upon Web objects. Finally, each pair of the encoding and decoding caching modules 270A, 280A, 270B, 280B, 270C, 280C can be a byte cache configured differently according to different fingerprint sizes. Each of

the encoding caching modules 270A, 270B, 270C can share the same memory space. Of import, a scheduler 300 can execute in memory of a host computer coupled to the server gateway 250. The scheduler 300 can include program code enabled to classify a request for data in the server farm 230 by an application in one of the clients 210.

The program code of the scheduler 300 further can be enabled to select one of the encoding caching modules 270A, 270B, 270C to perform data deduplication (or the proxy cache 270N, 280N) based upon the classification of the request for data. In this regard, in one aspect of the embodiment, the scheduler 300 can consult a table correlating different classifications with different ones of the encoding caching modules 270A, 270B, 270C. The classification can include, for example, an identity of an application associated with the request, a network address of the application, a port for communicating with the application, or any combination thereof.

The table itself can be constructed according to training data submitted to each of the encoding caching modules 270A, 270B, 270C from different servers according to different protocols. The performance metrics of each of the encoding caching modules 270A, 270B, 270C can be monitored and measured, for example, for throughput, processor and memory utilization and response time, to name only a few examples. The measured metrics can be submitted to a performance function, for example a function that weights different metrics for the utilization of different resources and then sums the weighted outcomes into an aggregated metric. The aggregate metric for a particular training data set can be compared against other aggregate metrics for the same training data set in different ones of the encoding caching modules 270A, 270B, 270C. In this way, the table can be constructed with an optimal one of the encoding caching modules 270A, 270B, 270C corresponding to a classification for the training data. Optionally, the performance of the encoding caching modules 270A, 270B, 270C can be monitored and fed back to the scheduler 300 so that the table can be adapted according to the monitored performance of the encoding caching modules 270A, 270B, 270C.

Once the scheduler 300 has selected a particular one of the encoding caching modules 270A, 270B, 270C, a byte stream of requested data by an application in a requesting one of the clients 210 can be deduplicated by the selected one of the encoding caching modules 270A, 270B, 270C and transmitted over a sideband channel to a corresponding one of the decoding caching modules 280A, 280B, 280C. The corresponding one of the decoding caching modules 280A, 280B, 280C can reconstitute the deduplicated byte stream and can transmit the reconstituted byte stream to the requesting application through gateway/firewall 260 by way of switch 240.

In even yet further illustration of the operation of the scheduler 300, FIG. 3 is a flow chart illustrating a process for dynamic caching module selection for optimized data deduplication. Beginning in block 310, a byte stream can be received from a server farm for delivery to a requesting application over a computer communications network. In decision block 320, it can be determined whether or not the requested byte stream already had been classified by the scheduler. If not, in block 330 the byte stream can be classified in accordance with a protocol utilized to transmit the request, a server address of a server transmitting the address and a port of the server associated with a requesting application. Thereafter, the classification can be returned to the scheduler in block 340.

In decision block 350, it can be determined whether or not the byte stream is to be deduplicated. In this regard, if the classification indicates that the protocol of the request does not lend itself to efficient deduplication, in block 360, caching can be bypassed and the process can end in block 390. Otherwise, in block 370 a particular caching module can be selected according to the classification and in block 380, the byte stream can be routed to the selected caching module. Finally, the process can end in block 390.

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electromagnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, radiofrequency, and the like, or any suitable combination of the foregoing. Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language and conventional procedural programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be

connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention have been described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. In this regard, the flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. For instance, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

It also will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

Finally, the terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a," "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations,

7

elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the principles of the invention and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

We claim:

1. A data deduplication data processing system configured for dynamic caching module selection for optimized data deduplication, the system comprising:

a server communicatively coupled to a data store and a plurality of client computers over a computer communications network;

middleware disposed between the server and the client computers and executing in memory of a host computer, the middleware comprising a plurality of different caching modules, each of the different caching modules having a different configuration and ones of the plurality of the different caching modules comprising program code enabled to deduplicate data in an identified particular one of the different caching modules; and,

a scheduler comprising program code executing in memory of a host computer and enabled to classify a request to retrieve data from the data store of the server, to identify from amongst the different caching modules a particular caching module associated with the classification of the request, and to route the data for

8

deduplication in the identified particular caching module, but bypassing the deduplication of the data when the request indicates that data is encrypted.

2. The system of claim 1, wherein the request is classified according to a table correlating different requests for different ones of the caching modules.

3. The system of claim 1, wherein the caching modules include byte caching modules each byte caching module configured with a different fingerprint size.

4. The system of claim 3, wherein the byte caching modules additionally include an object caching module.

5. The system of claim 4, wherein the object caching module compresses objects.

6. The system of claim 2, wherein the table correlates a protocol and network address for each of the different requests with a corresponding one of the caching modules.

7. The system of claim 1, wherein the program code of the scheduler bypasses deduplication of the data responsive to an indication in the scheduler to bypass deduplication of the data referenced by the request.

8. The system of claim 1, further comprising a table, wherein the table includes entries determined by processing training data for different requests in each of the different caching modules and correlates the request with an optimal one of the modules.

9. The system of claim 8, wherein the table is constructed by:

submitting the training data to each of the modules from different servers according to different protocols,

monitoring performance metrics of each of the modules and measuring the performance for throughput, processor and memory utilization and response time,

submitting the measured metrics to a performance function that weights different metrics for utilization of different resources and then sums the weighted metrics into an aggregated metric,

comparing the aggregate metric for a particular training data set against other aggregate metrics for the training data set in different ones of the modules, and

selecting an optimal one of the modules corresponding to a classification for the training data.

* * * * *